JOURNAL OF CYBERSECURITY

Journal of Cybersecurity, 2024, tyae023
https://doi.org/10.1093/cybsec/tyae023

Research Paper

A path forward: improving Internet routing security by enabling zones of trust

David Clark [1,*], Cecilia Testart[2], Matthew Luckie[3], kc claffy[3]

[1]CSAIL, MIT, 32 Vassar Street, Cambridge, MA 02139, United States
[2]Schools of Cybersecurity and Privacy, Georgia Tech, North Avenue, Atlanta, GA 30332, United States
[3]CAIDA, UCSD, 9500 Gilman Drive, CA 92093, United States

*Corresponding author. CSAIL, MIT, 32 Vassar Street, Cambridge, MA 02139, United States. E-mail: ddc@csail.mit.edu

Received 4 February 2024; revised 30 August 2024; accepted 10 October 2024

Abstract

Although Internet routing security best practices have recently seen auspicious increases in uptake, Internet Service Providers (ISPs) have limited incentives to deploy them. They are operationally complex and expensive to implement and provide little competitive advantage. The practices with significant uptake protect only against origin hijacks, leaving unresolved the more general threat of path hijacks. We propose a new approach to improved routing security that achieves four design goals: improved incentive alignment to implement best practices; protection against path hijacks; expanded scope of such protection to customers of those engaged in the practices; and reliance on existing capabilities rather than needing complex new software in every participating router. Our proposal leverages an existing coherent core of interconnected ISPs to create a *zone of trust*, a topological region that protects not only all networks in the region, but *all directly attached customers* of those networks. Customers benefit from choosing ISPs committed to the practices, and ISPs thus benefit from committing to the practices. We discuss the concept of a zone of trust as a new, more pragmatic approach to security that improves security in a region of the Internet, as opposed to striving for global deployment. We argue that the aspiration for global deployment is unrealistic, since the global Internet includes malicious actors. We compare our approach to other schemes and discuss how a related proposal, ASPA, could be used to increase the scope of protection our scheme achieves. We hope this proposal inspires discussion of how the industry can make practical, measurable progress against the threat of route hijacks in the short term by leveraging institutionalized cooperation rooted in transparency and accountability.

Keywords: Internet security; Border Gateway Protocol; BGP; trust zone

Introduction

The Internet's global routing protocol—Border Gateway Protocol (BGP)—suffers from a well-documented vulnerability: a network [termed an autonomous system (AS)] can falsely announce that it hosts or is on the path to a block of addresses (a *prefix*) that it does not in fact have the authority to announce. Routers that accept a forged route announcement—known as a *route hijack*—will route traffic intended for addresses in that block to the rogue AS. The simplest form of route hijack is an *origin hijack*, in which a malicious AS falsely announces ("originates an assertion") that it directly hosts (i.e. is the origin for) a prefix that belongs to someone else. In a *path hijack*, an attacker claims to be an AS *in* the path to a prefix, forging the legitimate owner's ASN as the origin of the prefix. The highly distributed operation of the BGP protocol—≈75K independent networks around the world—and its role in establishing and maintaining the connectivity we call "the Internet," have contributed to the persistence of this long-standing but increasingly dangerous vulnerability.

The two clear victims of a route hijack are the owner of the hijacked block and the sender of traffic to the hijacked block. If the attacker hijacks address space in order to impersonate the legitimate holder [1–4] or to inspect [5] the traffic, then senders of traffic to the hijacked block may fall victim to a scam or surveillance. If the attacker hijacks address space in order to conduct malicious activity

© The Author(s) 2024. Published by Oxford University Press. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs licence (https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits non-commercial reproduction and distribution of the work, in any medium, provided the original work is not altered or transformed in any way, and that the work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

[6–8], a third victim is the target of the malicious activity. The malicious activity may cause blocklisting of the address block, which impairs the legitimate owner's use of the block.

The best currently available practices in routing security require two steps to identify and block propagation of bogus route announcements. First, each ISP must register its own address space in a trusted database (ideally, the Resource Public Key Infrastructure aka RPKI), and routers across the Internet must check announcements against such a database and drop those announcements that are not consistent with the registered information [*route origin validation* (ROV)]. An AS who engages in the first step gains no security unless other ASes correctly deploy the second (ROV) step. However, ROV is sufficiently operationally complex that smaller or lower-resourced ISPs are reluctant to risk misconfigurations that impair their own service availability. Thus, early adopters of either technology *Route Origin Authorization*s (ROA or ROV) take on additional costs and operational risks, but the benefits may not accrue to them or their customers. Even if consistently implemented, which is a lofty aspiration in a global context, these two practices target only the simplest form of hijack, an *origin hijack*. A proposed approach for protection against a broader range of hijacks is BGPsec, an even more complex and expensive protocol-based solution that is at least a decade away from significant operational deployment [9].

Concerns over slow progress on routing security solutions led the US Federal Communications Commission (FCC) to issue a February 2022 Notice of Inquiry into potential regulatory interventions that could reduce the severity of the threat to US networks and traffic [10]. Several US government agencies, including the DHS and a joint filing by the DOD and DOJ, urged the FCC to take action [11,12]. Other commenters emphasized the risks of regulation in this domain.

Tension is increasing on this topic, as multistakeholder efforts to advance routing security have continued for over a decade. In the meantime, the risk and prevalence of both accidental and malicious BGP hijacks grows, rendering even the largest companies in the world victims of hijacks [3]. The scope of the problem is elusive to measure given lack of disclosure—and sometimes lack of awareness of—incidents.

The collective-action characteristic of the problem is fundamental: even those who are willing to invest in order to increase their own routing security cannot achieve protection without commitments from other networks to prevent propagation of bogus routes. We propose a more practical solution that refocuses on a new goal: to *provide a concrete action that a security-aware AS can take to protect itself from both having its address blocks hijacked, and its traffic to other address blocks hijacked*.

We propose an approach that achieves four related goals. First, it aligns incentives of actors toward improved routing security. Second, it offers protection against not only origin hijacks but the larger looming problem of path hijacks. Third, it allows ASes participating in the approach to protect their customers without additional work on the part of the customer, thus allowing highly resourced ISPs to protect other parts of the Internet. This feature is compelling because in today's Internet, the steps necessary to securely configure systems are sometimes complicated, and smaller ASes may not have the skills or resources to undertake them. Even more compelling is the resulting alignment of incentives: customers will prefer a participating provider since they offer enhanced security, giving providers an incentive to participate so they can market their improved security to potential customers. Finally, our approach requires no new capabilities in routers, relying on existing capabilities and institutions, and current techniques for analyzing interdomain (BGP) topology data.

The roadmap of this paper is as follows. We first provide some background and review related work, focusing on the barriers to routing security over the last two decades. We then dedicate a section to describing the threat model, including threat actors and defenders. We then introduce the principles of a *zone of trust*, a connected region of the Internet where providers take enhanced steps to improve the security of that region, including the security of customers connected to providers in the region. We introduce a specific example of a routing zone of trust that offers a more incentive-aligned direction for protecting ASes from both *origin* hijacks and *path* hijacks.

We analyze the residual risks of our scheme and how to minimize them [see section "Evaluating residual risk (local regions)"], auditing requirements (see section "Auditing requirements"), and comparison to other proposals (see section "Comparison to other proposed solutions").

Background and related work

The Internet standards community has long struggled with proposals to tighten the integrity of BGP communications. As with protection of other Internet transport mechanisms (e.g. DNSSEC, TLS), the standards community has grappled with complexities of cryptographic key management, trust anchors, and performance implications that hinder standardization, implementation, and deployment. Over the last 30 years, over 20 proposals to secure BGP have come from academia, industry, and the Internet Engineering Task Force (IETF), some of which Fig. 1 highlights. We describe how the standards and operational communities have tried to tackle this problem, and how it motivates our proposal.

Interdomain routing

ASes use BGP to exchange *routes* that describe paths to destinations in the global Internet. Two important components of a route are the *prefix* that specifies the block of addresses of a route, and the *AS path* that reports the sequence of ASes that received the route. Internet traffic will flow back to the destination prefix following the AS path. To prevent forwarding loops, a router chooses the most specific route to a destination IP address—i.e. for 192.0.31.8, it would prefer a route with a prefix 192.0.31.0/24 over 192.0.30.0/23. Operators use this property for traffic engineering. BGP also provides a mechanism to annotate announcements with attributes—known as *BGP communities* [13]—to enable signaling within and across ASes, facilitating traffic engineering innovations [14] such as automated blocking of denial-of-service attack traffic on the path to the victim [15,16].

There are two general types of relationships between neighboring ASes: customer-to-provider (c2p), where the customer pays a provider to obtain global reachability, and peer-to-peer (p2p), where two peers exchange routes to their customers without involving an intermediate provider [17]. If an AS has multiple routes to the same prefix, the rational choice is to prefer routes received from customers (a source of revenue), over routes received from peers (typically settlement-free, i.e. no cost), over routes received from providers (which cost the AS) [17]. Other ASes that an AS X can reach through a customer link are within the *customer cone* of X.

A few (≈15) large ASes obtain global routing using routes received only from their peers and customers, i.e. they do not pay any transit providers. These ASes connect in a full mesh (a peering clique) that enables packet delivery between arbitrary networks with different transit providers. The ASes in this group that also do not pay for peering are known as Tier-1 providers. However, payments between ASes are confidential; we thus use the term *Tier-1* to refer to the peering clique.

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

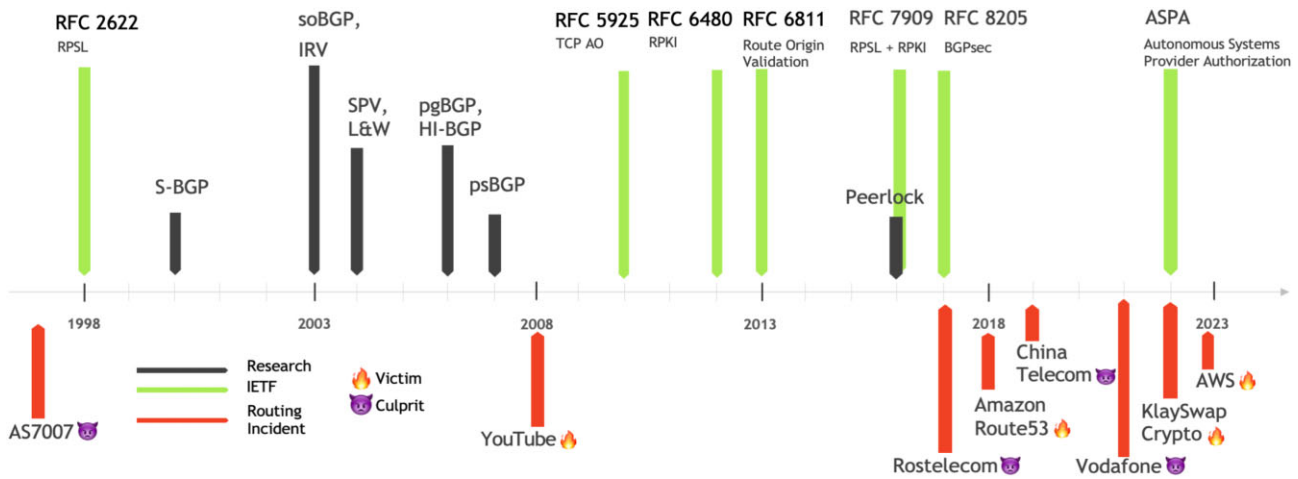

Figure 1. Decades of proposed routing security approaches; sample of high-profile hijacks.

Routing security in the 1980s

In 1982, Rosen [18] documented that it is possible to corrupt inter-domain routing in RFC 827, in the context of a predecessor of BGP called the Exterior Gateway Protocol (EGP):

> If any gateway sends an NR [neighbor reachability] message with false information, claiming to be an appropriate first hop to a network which it in fact cannot even reach, traffic destined to that network may never be delivered. Implementers must bear this in mind.

This warning to implementers suggests that the perceived threat in 1982 was accidental misconfiguration, rather than malicious operators.

Routing security in the 1990s

To mitigate the prevalent risk of accidental misconfigurations, in the 1990s network operators developed the *Internet Routing Registry* (IRR) system of distributed databases. The IRR system enabled network operators to publish address ownership and routing policy information [19], which other operators could use to build filters that permit or deny routes according to these operator-registered policies. Unfortunately, some IRR databases do not fully authenticate registration data, allowing attackers to compromise the IRR by falsely registering ownership of resources, which they then use in a hijack [20–25].

When the IETF started to study malicious BGP security threats in the late 1990s, they did not initially assume that an AS operator was an important threat actor. Instead, they focused on the threat that a third party could intercept the traffic between two well-behaved ASes and then modify the BGP update to inject a false assertion. To defend against this threat, in 1998 the IETF added an optional extension to TCP to allow end-points to authenticate the contents of a TCP segment [26,27].

Routing security in the 2000s: proposals for a secure BGP

Aiming for a more complete approach to routing security, in 2006 the IETF's Secure Inter-Domain Routing (SIDR) Working Group began designing a variant of BGP that would support *path validation*—ensuring that each AS appearing in a received AS path was legitimately in the path. During this decade, over a dozen competing approaches came out of academia and industry [28–41]. The protocol that became an IETF standard (RFC 8205) in 2017 is called BGPsec [42]. BGPsec update messages include two important new fields: the AS to which the router is sending that announcement; and a cryptographic signature over the message that enables any router along the path to verify that the series of signatures are valid. This mechanism prevents path hijacks: a malicious AS cannot forge the AS path because the malicious AS cannot sign records for the forged ASes.

Cryptographic attestation of paths requires propagation of a new layer of cryptographic transaction at each hop, which is computationally expensive and poses a router-level (rather than AS-level or prefix-level) key distribution challenge, since every router must have its own public key signed by a certificate authority. Furthermore, full protection of the path requires every AS along the path to implement BGPsec. Partial deployment, inevitable during a transition, implies unpredictable protection. The complexity, overhead, and misaligned incentives have prevented significant operational deployment of BGPsec, despite a decade-long standardization process that completed in 2017 [42].

Routing security in the 2010s

The 2010s brought three areas of endeavor: rigorous analyses of the incentives to deploy routing security solutions; technology, standardization, and operational mechanisms to mitigate the simpler problem of origin hijacks; and a collective action effort termed Mutually Agreed Norms for Routing Security (MANRS) that aimed to overcome the counter-incentives to deploying these mechanisms.

Analyzing deployment incentives

As early as 2009, researchers began to survey the array of routing security efforts and analyze why they had failed to gain traction [43,44]. Such reviews continued throughout the subsequent decade [45–47]. Researchers also explored approaches to overcome the economic counter-incentives to deployment of protocol-based approaches to routing security, and analyzed the implications of partial deployment of such approaches [48–51]. The deepest body of work on this topic was by Sharon Goldberg and Michael Schapira and their collaborators.

In 2011, Gill, Schapira, and Goldberg proposed a strategy that would create market pressure to adopt BGP path validation. (They referred to the set of options at the time as *S*BGP*). Their proposal

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

required (e.g. by regulation) a few Tier 1 ISPs to first deploy S*BGP, and required those participating in S*BGP to prefer secure routes over other routes to the same prefix [48]. This scheme also reduced deployment complexity by allowing transit providers to cryptographically sign routes on behalf of their stub customers. Their simulations on realistic AS topologies showed that under these conditions, the S*BGP ASes would draw traffic away from other ASes, and most of the rest of ASes would then switch to S*BGP to get their traffic (revenue) back. Followup work two years later [49,50] acknowledged that having Tier 1 ISPs lead a market-driven deployment would not work because economic incentive would override any secure route received from a peer when an insecure route via a customer was available.

In 2011, researchers proposed a new Internet architecture, SCION [52], which separated ASes into independent trust domains that provide isolation of routing failures and misconfigurations. SCION assumes a hierarchical architecture, where one or more highly trusted ASes connect the domains to each other. Researchers recently described how SCION could be used to bootstrap a secure routing system [53], but it has not achieved significant interdomain use [54]. In contrast to the approach we propose here, which builds on existing protocols, SCION is what has been called a "clean slate" design, including the replacement of BGP with a new protocol and changes to the Internet Protocol (IP) header.

In 2016, Cohen *et al.* [51] proposed an approach similar to the subsequently proposed Autonomous System Provider Authorization (ASPA) protocol (which we describe later in this section). Their simulations focused on the length of paths that an attacker must construct if the AS announcing the prefix has registered what we today call an ASPA. The authors discussed deployments in select geographic regions, perhaps driven by government pressure. They did not propose a connected region, so partial deployment of this approach yields only a probabilistic assessment of protection, as with ROV (and ASPA).

Preventing origin hijacks: RPKI and ROV

While BGPsec has been undergoing implementation and evaluation for a decade, operators have focused on the more tractable challenge of ROV, which is recognized as the best current practice in routing security. The IETF SIDR WG specified ROV in 2013 as a mechanism to mitigate the risk of *origin hijacks* (the simplest form of hijack) [55]. ROV uses a Resource Public Key Infrastructure (RPKI) [56], an authoritative database maintained outside of BGP that closely matches Internet resource delegation by the Regional Internet Registries (RIRs). Each of the five RIRs is the root of trust, enabling the holders of the IP addresses blocks delegated by them to issue ROAs. ROAs are cryptographic signatures that authorize designated ASes to originate routes to address blocks. Routers using ROV drop BGP announcements that are not consistent with a registered ROA for the prefix. RFC 6811 [55] specifies the ROV protocol with important caveats: its dependence on the integrity of the database used to validate routes, and its inability to prevent path hijacks. This residual risk includes the forged-origin path hijack mentioned above, where the malicious AS impersonates the valid source AS by appending it to a forged BGP announcement (recently observed in the wild [25]). RFC 6811 thus cautioned: ".. this system should be thought of more as a protection against misconfiguration than as true 'security' in the strong sense."

Use of ROAs presents other operational challenges. A ROA contains a set of prefixes and a set of origin ASNs. For a BGP route to be valid according to RPKI, there needs to be a ROA with a prefix that is equal to or covers the route prefix, the ASN originating the BGP route has to be in the allowed set, and the length of prefix needs to be allowed. Indeed, ROAs may also contain a *maxLength* attribute that defines the maximum prefix length allowed for each prefix; for example, an ROA for 192.0.30.0/23 with a maxLength of 24 enables the AS to originate 192.0.31.0/24. Operators use this feature for traffic engineering (see the discussion of Interdomain Routing in section Background and Related Work). In 2017, Gilad *et al.* showed that use of the maxLength attribute could enable an attacker to hijack more-specific prefixes that victim networks then unwittingly communicate with. Best current practice is to not use the maxLength attribute [57]. Furthermore, if a BGP route has a prefix covered by a ROA but the route is not valid either because the origin AS is not in the allowed set or the prefix length is not allowed, the route is invalid according to RPKI. As a consequence, if a network provider registers a ROA for a large prefix (e.g. a/16), any sub-delegation to another (smaller) network will be covered by the ROA, and more specific routes to those smaller prefixes may be considered invalid. RPKI requires coordination between network operators to prevent making routes unreachable.

Although RIRs have supported RPKI registration of ROAs since 2013, until 2019, there was little evidence of Internet Service Providers (ISPs) using ROAs to validate BGP announcements. However, by late 2022, many large ISPs, including AT&T, KPN, Arelion, and Comcast had started to use ROV to drop invalid announcements [58–61]. According to NIST's public RPKI monitor based on RouteViews data [62], as of May 2024, 51% of IPv4/24s in unique prefix-origin pairs advertised in BGP were covered by RPKI and observed as valid, i.e. the origin AS in the BGP announcement matched the registered ROA. These statistics vary by region: for May 24, 2024, NIST reported 70% of observed prefix-origin pairs in the RIPE region were valid, 58% in LACNIC, 51% for APNIC, 38% for ARIN, and 30% in the AFRINIC region [63]. Recent work examining the state of ROV deployment in the Internet between December 2021 and September 2023 reported that 12.3% of tested ASes had behavior suggesting that they or all of their transit providers had consistently implemented ROV [64]. They reported that larger ASes (i.e. those networks with technical capacity) were more likely to have implemented ROV, consistent with other studies [65]. APNIC has a live ROV measurement based on the availability of a path to a prefix that switches RPKI status every 2 to 3 days [66], and reports results at a per-country level in a world map [67]. In May 2024, the world map showed the disparities in ROV adoption at that time. Many countries, including the USA, France, Spain, Sweden, Finland, and Australia had over 50% adoption. Others such as Russia, China, Brazil, and Mexico still had less than 10% ROV adoption. Using BGP and RPKI data to infer when networks drop invalid origin announcements following the methodology from [68], we found that about 60% of the 360 networks sharing their data with BGP collectors had adopted ROV. Tier-1 and large networks in the USA and Europe are more likely to share BGP data with collectors. 14 out of 17 Tier-1 networks measured with that methodology have deployed ROV.

Collective action attempt: MANRS

In 2014, several network operators established a voluntary initiative to promote operational practices to "help reduce the most common routing threats on the Internet," which they called *Mutually Agreed Norms for Routing Security* [69]. MANRS specified four practices for participating networks, two of which correspond to the RPKI/ROV steps of registering authoritative information about one's prefixes, and verifying BGP announcements against authoritative information. The exact wording of these two practices is as follows: (1) *prevent propagation of illegitimate routes from customer networks or one's*

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

own network. and (2) *document in a public routing registry the prefixes that the AS will originate*.

To conform with the first practice, a MANRS member must verify two aspects of an announcement from a customer: (1) it must confirm that the customer has used an ASN that it is legitimately allowed to use and (2) for any prefix originated by that customer, that the ASN is allowed to announce that prefix. However, to encourage broad uptake, MANRS does not specify how a member AS should verify the assertions of its customers, and in particular does not require the use of RPKI/ROV (ROAs) in this verification. The AS can use ROAs, or can verify against (less authoritative) information in the Internet Routing Registry (IRR), or rely on a private arrangement with its customer.

The MANRS initiative has a key strength: it illustrates that ISPs can institutionalize their recognition of the need for a collective commitment to operational practices to reduce threats to the routing system. However, as the US Federal Communications Commission (FCC) observed [10], the MANRS program has had limited success. In May 2024, MANRS had 938 ISP and 30 CDN organizational members that operated 1268 and 30 ASNs, respectively [70]. This constitutes 1.7% of the ≈75K routed ASes. Many of the largest ISPs do not participate in MANRS, and some participating ISPs are not conforming to the MANRS practices. Du *et al.* reported their analysis of data from May 2022, at which time 5% of MANRS ISPs did not conform with the requirement to register their prefixes in either RPKI or IRR and 16% did not conform with the filtering requirement [71].

The limited success of MANRS (and its underlying practices) is rooted in misaligned incentives that manifest in three ways. First, although if consistently implemented, the MANRS practices will reduce the incidence of invalid origin hijacks, *there is no direct relationship between the action of any given MANRS member and the overall security of the Internet, or even the security of any customer of a MANRS member.*

Second, the current MANRS practices, even the stronger RPKI/ROV options, only aim to prevent origin hijacks rather than path hijacks. Some network operators believe this benefit does not justify the cost and complexity of RPKI/ROV.

Third, there is insufficient auditing of conformance to lend confidence in assuming consistent implementation [72]. Independent auditing has detected significant non-conformance to MANRS practices [71,73]. More rigorous auditing would be expensive and further reduce the incentive to participate.

Autonomous System Provider Authorization

Recognizing the barriers to BGPsec deployment, and the lack of path validation capability in ROV, in 2019, several engineers proposed Autonomous System Provider Authorization (ASPA) as a mechanism to protect against route leaks and forged-origin prefix hijacks [74]. As of June 2024, ASPA is still in IETF development. ASPA builds on presumed use of RPKI and ROV but enables customer ASes to go further by registering a list of their transit providers in the globally visible RPKI database. That database allows any AS to examine a BGP announcement to detect and reject many types of invalid path announcements, so long as the ASes along the path have registered their providers in ASPA. The authors describe ASPA as preventing route leaks as well as some forms of path hijacks; ASPA does not prevent an attacker from spoofing a sequence of ASes in the path if those ASes do not implement ASPA. Our proposed scheme provides a more predictable level of protection and improves incentives for deployment. We later compare our proposed scheme to ASPA, and describe a way in which the use of ASPA could expand the protection provided by both approaches.

Routing security in the 2020s

This decade, routing security has caught the attention of regulators. Researchers discovered hijacks of unannounced address space [75], and forged-origin hijacks of RPKI-valid address space [25]. After earlier hijacks of AWS address space [2] motivated Amazon to register ROAs for most of its address blocks, attackers developed more sophisticated path hijacking techniques. The high-profile hijack of AWS space in August 2022 [3] motivated by the opportunity to steal cryptocurrency,

succeeded for multiple reasons. Amazon signed multiple ROAs that allowed different ASNs to originate their prefix; these ROAs had *maxLength* attributes that the attacker exploited to hijack an IPv4/24 that hosted the crypto-currency service; and the attacker registered that IPv4/24 in an unauthenticated IRR entry to convince upstream providers to permit the prefix announcement. However, even if Amazon had announced a competing more specific prefix, the attacker's path would have been preferred for networks that were customers of AS1299 who did not have a more-preferred route to Amazon.

The persistent failure of market-driven solutions to routing security has recently triggered government interest and inquiry into potential interventions. In 2022, the OECD [76], ICANN [77], BITAG [78], and the US FCC [10] all published reports with extensive references related to routing security challenges, and limitations of proposed solutions.

We expect governments to feel compelled to intervene in the Internet infrastructure ecosystem to improve routing security, and we seek to provide an alternative that leaves as much control as possible with the participating networks. Our approach draws inspiration from Lychev *et al.*'s conclusion a decade ago [49] regarding market-driven evolution of secure routing: "We hope that our work will call attention to the challenges that arise during partial deployment, and drive the development of solutions that can help surmount them.. Alternatively, one could find deployment scenarios that create 'islands' of secure ASes that agree to prioritize security 1st for routes between ASes in the island; the challenge is to do this without disrupting existing traffic engineering or business arrangements" [49].

We pursue this challenge with an approach that leverages a coherent topological region to achieve our design goals: incentive alignment, competitive advantage to participating networks; proportional responsibility, in that larger players can invest to protect their customers, providing this competitive advantage; and protection against origin as well as path hijacks without the operational complexity of BGPsec. We believe our proposed alternative is worth open debate before pursuing more blunt regulatory measures.

Threat model

We next describe the capabilities of defenders, to contrast defender capabilities with attacker capabilities.

Defender capabilities

As of May 2024, ≈85% of the ≈75K ASes on the Internet have no customers. They are in many cases small ASes with limited operational resources to defend themselves. While they may use peering connections to handle some of their traffic [see section "Evaluating residual risk (local regions)"], they connect to transit providers to reach most parts of the Internet. These transit networks engage in contractual agreements when they interconnect with their neighbors. These transit network operators regularly interact at peering forums and other industry events (e.g. NANOG) and thus have established relationships. In our threat model, the defenders are these transit

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

providers. Defenders have the capability to establish parameters with their customers in terms of what prefix announcements the customer is expected (and allowed) to make, and thus to automatically accept or reject routes through configuration capabilities present on routers. Defenders can access external databases, e.g. IRR, RPKI, to support their assessment of their customer routes.

A defender does not in general have the ability to verify the announcements of their customers' customers, due to the temporal dynamism in the interdomain relationships of their customers. Further, some defenders and their customers are limited in how they use RPKI. For example, some legacy resource holders are hesitant to obtain ROAs, as doing so would require they enter a contractual agreement with an RIR. (In particular, ARIN's agreement embedded a controversial position that in order to register a ROA, holders of legacy space (i.e. address blocks allocated before ARIN existed) must contractually agree that they have no legal property rights to their address space [79]. In September 2022, ARIN removed this clause from their Registry Services Agreement, but asserted that ARIN still did not recognize property rights in IP addresses [80].) Finally, a defender cannot control the route selection policies of their peers or customers; these ASes might select hijacked routes from other neighbors they have.

Attacker capabilities

We assume that the attacker controls or has subverted an AS that connects to the Internet using one or more transit providers, which provide routing to the rest of the Internet for that AS and deliver traffic intended for that AS. The attacker has the ability to corrupt *unauthenticated* databases, such as IRRs, with false claims that the attacker is the legitimate holder of a prefix. Finally, an attacker has the ability to commit to security practices that they have no intention to follow.

An attacker does not have the ability to completely hide their activities; in order for their attack to be effective, their hijacked route must propagate, and multiple route collector projects today publish the set of AS paths they collect. Nor does an attacker have the ability to issue ROAs for address space that they do not control. An attacker could compromise the RIR (an insider) or the prefix holder's RIR account, but that is out of scope for the proposed approach as it is a generic and well-understood security problem of systems connected to the Internet.

A routing zone of trust

We first introduce the concept of a *zone of trust* in a routing context, before specifying a zone of trust in more detail below in our discussion of the Verified IP zone (VIPzone). Figure 2 depicts a zone with member providers (in green) at the edge of the zone providing transit service to directly attached customers (white). The providers connect within the zone, and must know when they are exchanging traffic with another member of the zone, and when they are communicating with an AS outside the zone.

A zone could protect against origin hijacks as follows. If all providers P in the zone commit to implement ROV and drop invalid origin announcements from customers outside the zone, then no invalid origin announcements will circulate inside the zone, which means that customers C will never receive a BGP announcement from the zone where the origin is invalid based on a ROA. These practices turn this zone into a *zone of trust*.

This example illustrates three properties of a zone of trust:

(i) *Collective action* by ASes creates the zone and its trust attributes.

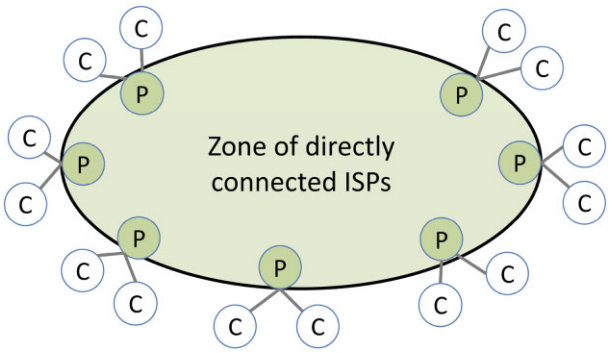

Figure 2. A requirement of our proposed *zone of trust* is a coherent topological region with providers in the zone providing transit to attached customers.

(ii) ASes in the zone must have paths that connect them to all other ASes within the zone, i.e. the zone is not partitioned.
(iii) Customers of providers in the zone obtain protection from invalid origin hijacks by using a provider in the zone. They need take no other action.

We call this region a *zone of trust* because the protection in the zone arises from actions of ASes at the perimeter of the zone. This protection requires that ASes in the zone be able to trust that the routers at the perimeter function correctly, which requires some degree of transparency and accountability. We introduce this design assumption in exchange for one that routing security protocols have always included: global deployment of a protocol. If ASes themselves can be threat actors, we are skeptical of an aspiration to make BGP *globally* secure. Creating a zone of trust through perimeter protection (a trust-but-verify regime) offers a more pragmatic approach for today's routing system.

Given the history of routing in the Internet, where each Autonomous System can choose to interconnect based on its own needs, the idea of a coherent perimeter around a zone is missing from today's interdomain routing system. ROV deployment discussions today consider each AS in isolation, leaving security a statistical measure. We can count the number of ASes that register their ROAs, or the number of ASes that implement ROV, but the consequence for a given AS is a function of what other ASes choose to do. It is thus not clear what specific action an AS should take to reduce its own risk profile. Today, invalid origin announcements may propagate across the Internet, and may or may not reach any given AS. In contrast, a connected zone of trust allows clear articulation of the benefit to a given AS to joining the zone: ASes in the zone will receive no announcements from the zone with an invalid origin based on a registered ROA.

The incentive alignment extends beyond the zone: customers concerned about hijacks can seek out providers that are in the zone, which in turn creates an incentive for providers to commit to the required practices that define the zone and join it. Today, there is little direct benefit to an AS that chooses to implement ROV. Many of the larger ASes do so, as part of a collective action to improve security, but recognizing that these actions can create a coherent zone with direct benefit to their customers will increase their incentive.

Note that the zone does not provide absolute protection from origin hijacks for ASes with connectivity outside of the zone. If a customer C has its own customers, peers, or other providers not in the zone, it could still receive a hijack from those nearby ASes. We call this set of ASes the *local region* of the customer C, and we characterize this residual risk in the section on Evaluating Residual Risk. Importantly, the residual risk depends on the size and character of

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

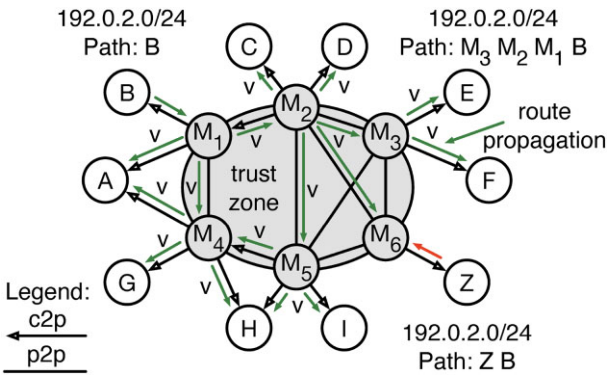

Figure 3. A Routing Zone of Trust can defend members and their customers from path hijacks in the zone if members (M) mark routes from their customers as VERIFIED(v) as they enter the zone, and other zone members select VERIFIED routes over unverified routes. Above, M_1 expects its direct customer B to announce 192.0.2.0/24, so M_1 marks that route as VERIFIED and propagates it to other zone members. Black arrows show c2p links, lines without arrows show p2p links, and green arrows show route propagation from B. The hijacked route via Z does not propagate in the zone because Z is not a member, and the zone has an alternative VERIFIED route.

the local region of each AS, which that AS can know and control according to their own risk profile.

Figure 3 allows us to consider how the trust zone gives each AS control over the two types of hijack harm: having one's addresses hijacked or having one's traffic hijacked. An AS (e.g. B) can protect against hijacking of its own addresses (which we call *owner harm*) in the zone by directly connecting to the zone, and registering its addresses in the RPKI. Other ASes that attach to the zone are thus protected from hijacking of their traffic to B's addresses (which we call *misdirection harm*). An AS that does not consider owner harm a significant risk need not register its addresses in a database (although we encourage universal use of the RPKI). The AS may care more about misdirection harm and might thus minimize its local region and get as many route announcements as possible from providers in the zone. Different ASes may have different risk assessments, and unlike today's routing ecosystem, this trust zone is structured to allow an AS to pick its own options based on its own assessment of route hijack risk.

Does a coherent topological zone already exist in today's Internet?

Could such a coherent topological region exist? In fact, it already does, in the context of the MANRS initiative. Many of the MANRS members make up a connected region today. In May 2024, MANRS had 938 ISP members, with 1268 ASNs [70]. To derive the connected region, we perform a topology exploration using CAIDA's ASrank data [81] for May 2024. We start with members with no providers (Tier 1 providers), and recursively add directly connected customers that are also MANRS members. The resulting region has 581 members with 766 ASNs. Currently, 28 592 AS-level customers directly connect to this region. If MANRS could extend their operational practices to make this region a zone of trust, more than one-third of the ASes active on the Internet today would receive that protection.

VIPzone: using verification tags to prevent path hijacks

We now describe how a set of proposed operational practices in a coherent zone of trust, which we call *VIPzone* (for Verified IP zone), will limit *path hijacks*. For an AS to be in the VIPzone, it must commit to these practices and must be part of a connected zone. To be part of the connected zone, it must either be a Tier-1 provider, or have a member of the VIPzone as a transit provider.

Figure 3 illustrates the basic VIPzone operation. We describe the VIPzone practices below, and provide a finer-grained specification of these practices in Appendix 1.

Our **VIPzone** builds directly on the current requirements of MANRS [69]. Today, each MANRS member is required to verify all announcements originated by its directly connected customers. The member must perform two checks: (1) that the customer has used an ASN that it is legitimately allowed to use and (2) for any prefix originated by that customer, the ASN is allowed to announce that prefix. The member must rely on direct knowledge of its customer (a "know your customer" or KYC requirement) to verify that the AS used is legitimate. We emphasize that the requirement that a provider verify the AS numbers that a customer is authorized to use can be difficult to implement, and attackers have subverted routing by exploiting this difficulty. This challenge is the same for MANRS and our VIPzone. The difference is that we propose an approach that shifts the *incentive structure* toward careful implementation—ISPs benefit from their own effort. BGPsec (see section Comparison to Other Proposed Solutions), which has not achieved any material deployment, is the only technical mechanism that can confirm (so long as the cryptography is not compromised) that a customer controls the AS that it is using for its announcements. A (MANRS or VIPzone) member can use RPKI validation, an authenticated IRR database, or a manually configured prefix list (ACL) to verify the non-member's announced prefix is correct. In our VIPzone scheme, the zone member then either drops such announcements or marks them VERIFIED. We propose the use of a BGP community value [13] to carry the VERIFIED marking, similar to a recent IETF proposal to use a community value to annotate path properties in order to allow detection of route leaks [82]. For announcements that come from customers of the customer not in the zone, the VIPzone member forwards them without marking them VERIFIED. Neither MANRS nor VIPzone requires that an AS check the validity of the path in an announcement with more than one AS in the path, due to the assumed infeasibility of this check.

However, for announcements that VIPzone members have verified, they must propagate the VERIFIED marking as they forward announcements within the zone. A member must remove this marking if it appears in any announcement entering from outside the zone. This allows VIPzone members to establish the authenticity of VERIFIED announcements, regardless of their distance from the origin. Finally, inside the zone, any AS receiving multiple announced routes for the same prefix must prefer one marked VERIFIED. By this rule, no member will prefer a path hijack route over a legitimate route from customers directly attached to the zone, since legitimate routes will be marked VERIFIED.

Customers directly connected to the zone minimize owner harm, both for origin and path hijacks. Zone members verify prefixes received from attached (non-zone) customers, and mark them VERIFIED before forwarding them into the zone. If a malicious AS directly connected to the zone tries to launch an invalid origin hijack, zone members will discard it based on the KYC practices. If the AS launches a path hijack (which must by definition have more than one AS in the path), the member AS may forward it unverified into the zone (a "not sure" situation), but it will have no impact so long as a corresponding VERIFIED announcement is active.

We emphasize the essential role of the VERIFIED tag. When a MANRS member cannot verify whether the path announcement is valid (e.g. multiple ASes in path), the member can forward this an-

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

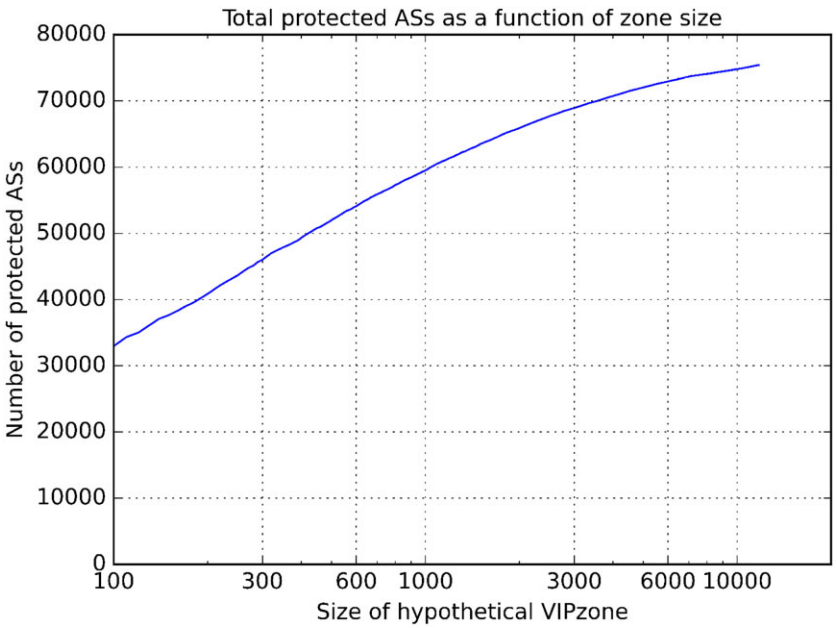

Figure 4. Protected ASes (in the zone or connected directly to it) as a function of zone size (ASRank data, May 2024).

nouncement onward. Forwarding potentially invalid announcements without *any* signal of risk prevents the current MANRS framework from manifesting a zone of trust. A key requirement of the VIPzone approach is that members propagate two sorts of announcements in the zone: VERIFIED and "not sure." This feature allows for more flexible and incremental deployment of the protections. In our VIPzone proposal, each AS drops invalid announcements, marks announcements as VERIFIED if it knows they are correct, and forwards announcements without the VERIFIEDmarking if the AS is not sure. The rule that makes the zone trustworthy in this case is that if there is a VERIFIED announcement for a particular prefix, and one that is not VERIFIED (e.g. not sure) for the *same prefix*, zone members must prefer the VERIFIED announcement. This rule constrains the routing policies of zone members, the implications of which we discuss in the section "Evaluating residual risk" and elaborate on more complex scenarios in Appendix 2.

Protection against route leaks

A route leak is an event in which an AS inappropriately (i.e. violating routing policy) forwards a route it legitimately received. The consequence is often that large flows of traffic reach this AS, which is not provisioned to carry them. A classic route leak occurs when a multi-homed AS that takes the routes it receives from one of its transit providers and inadvertently propagates these routes to its other transit provider.

In addition to preventing path hijacks of ASes directly attached to the zone, the VIPzone prevents leaks of announcements of prefixes by ASes *not* in the VIPzone. If the leak occurs within the zone, the announcement would be VERIFIED and thus propagated within the zone. This potential harm from accidental misconfiguration suggests an important insight: most ASes should not be in the VIPzone, but should get the protections by being a *customer* of a VIPzone member. We consider it preferable that only operators with sufficient technical abilities and resources join the VIPzone. We elaborate on this idea in Appendix 3.

Protection against sub-prefix hijacks

One hijack that can penetrate the zone is based on a sub-prefix (an address block that is a subset of a VERIFIED prefix). Normal routing rules require that an AS, when selecting among routes for an arriving packet, must prefer the announcement with the longer prefix (i.e. smaller address block). Note that requiring that a VERIFIED announcement for a given prefix take precedence over an unVERIFIED announcement for a longer prefix risks breaking traffic management practices that disaggregate prefixes. Such a requirement could introduce loops. An AS concerned about owner harm resulting from a sub-prefix attack protects itself by registering ROAs for the prefix.

Evaluating VIPzone protections

We explore how many ASes would receive protection from hypothetical zones based on today's Internet topology. Using CAIDA's AS Rank data from May 2024, we initialize a zone with the 100 ASes with the largest customer cones. We then add new members, again ordered by the size of their customer cone. Figure 4 shows the number of protected ASes expanding rapidly with zone size, up to 11 781 ASes (in this data set), at which point every AS with any customers is in the zone. The only ASes not in the zone are single AS stubs.

However, note that such a large zone is unrealistic. Most ASes in that zone are small providers with few customers, likely without sufficient operational sophistication or resources to join the zone. If we pick an arbitrary cutoff of 600 members (about the size of the current MANRS zone), that would protect a little over two-thirds of the ASes in the Internet (in this hypothetical analysis, 53 563). This number is higher than the 28 592 customers of the current MANRS region we discussed earlier, because this VIPzone is formed by including all of the largest ASes (even non-MANRS members) as measured by their customer cone.

Social engineering attacks

As protection against traditional hijacks improves, attackers devise new ways to disrupt routing. One is a social engineering attack in which an attacker contacts a provider of a target AS, and (pretending to be an agent of the target AS) requests that the provider provision a new link to serve that target AS. If the provider does not recognize that the request is not legitimate, the attacker now has a

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

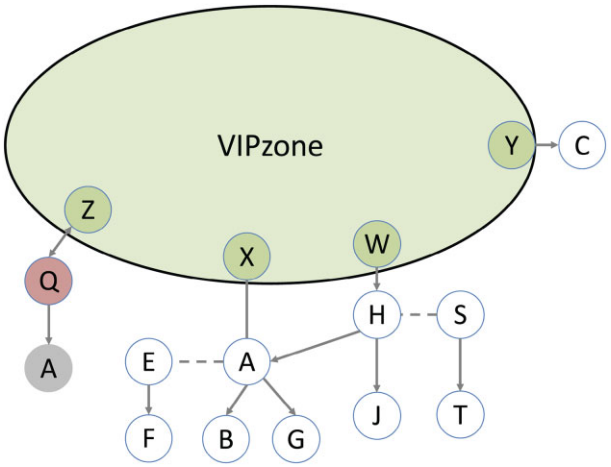

Figure 5. Various customers of a VIPzone, including A with a local region, C with no local region, and a malicious AS Q pretending that A is a customer.

BGP connection to the provider that the provider thinks is associated with the target AS. At this point, the attacker can announce routes (e.g. hijack them) associated with the target AS, and the provider will accept these announcements.

Transit providers will have to harden their implementation of the MANRS Know Your Customer requirement to detect these sorts of attacks. If the attacker can bypass the KYC test via impersonation of the address owner, ROAs and ROV are of no use in prevention, since, in that context, it appears a correct AS number is being used. This requirement applies equally to the existing MANRS, our VIPzone, and ASPA.

Evaluating residual risk (local regions)

We review the residual risks that an AS faces even if it is directly connected to the VIPzone, and what that AS can do to reduce these risks. We have already described how an AS mitigates the risk of *owner harm* simply by connecting to the zone. The residual risk of *misdirection harm* if they connect to the zone is a function of the size of the *local region* and the probability that a malicious AS operates in that region.

A local region of a VIPzone customer AS arises due to that AS's interconnection arrangements outside the zone, from which the AS receives BGP announcements. These include other ASes in the customer cone of that AS, the peers of that AS and their customer cones, and any providers (and their neighbors, recursively) of that AS that are not in the zone. In Fig. 5, A has provider X in the zone. Its local region includes customers B and G, peer E and E's customer F, provider H (which is not in the zone), H's customer J, peer S of provider H, and S's customer T. If provider H itself had a provider that was not in the zone, that provider, its customers, and any peers of that provider and customers of those peers would also be in the local region of A. Any of these could launch a hijack that triggers a *misdirection harm* to A.

We make three observations about local regions. First, the risk of hijack by one's own customers (A's customers B and G in this example) is a function of the risk of malicious behavior in the local region. But A (or any AS outside the zone) can mitigate this risk by implementing a robust KYC practice, which can generally detect forged-origin attacks by customers.

Second, misdirection from a hijack in the region is restricted to the region. In Fig. 5, if malicious AS Q launches a path hijack asserting that it has A as a customer, that announcement may penetrate the zone but without a VERIFIED mark, so zone members will prefer the VERIFIED announcement from X.

Third, for many attached customers, the local region is small. To examine the size distribution of local regions, we return to our hypothetical VIPzone (i.e. seeded with 100 ASes with the largest customer cones) and compute the size of the local regions for all attached customers. We add to the zone 100 ASes at a time, and at each step compute the size of the local region for the attached customers.

Figure 6 plots the resulting distribution. To compensate for the limited observability of peering relationships in public BGP data [83], we use two methods to compute the size of the local region. Figure 6a plots the local region size using the customer-provider and peering relationships from CAIDA's ASRank data [84]. Figure 6b relies on the method in [49], which augments observable peering relationships by assuming that any two ASes that attach to the same IX have a peering relationship. We use data from PeeringDB and PCH [85] to augment the set of peering relationships inferred by AS Rank.

Figure 6a underestimates the sizes of local regions, since CAIDA's ASRank data is derived from BGP announcements collected by RouteViews and RIPE RIS [86,87], and those vantage points do not have sufficient density to capture all peering relationshops. Figure 6b probably overestimates the sizes of the local regions, since many ASes that connect to IXs have selective peering policies. So the actual distribution probably lies between these two sets of curves. Current publicly available data does not allow a finer-grained estimate.

Note that the distribution of local region sizes in our data set is bimodal. Depending on the zone size, between 30% and 60% of the customer ASes have a very small local region—close to 1 AS. These are stub ASes that obtain access to the Internet using a transit provider, and do not peer to obtain connectivity. The right side of the plots shows large local regions, which represent ASes that peer widely to reduce their dependency on transit providers, or else use multiple providers, one of which is not in the zone, and which itself uses massive peering. A realistic consequence of extensive peering with ASes that

do not take known steps to verify their announcements is an increased risk of hijack. That expanded attack surface in the ecosystem is a motivation for the approach we propose.

Why we cannot realistically assess residual risk using current topology data

This analysis provides a hypothetical indication of the level of protection and residual risk that a VIPzone would yield under current interconnection patterns. But it is a problematic approach to assessing residual risk, since the architecture of the VIPzone will affect peering incentives by design. That is, the goal of VIPzone proposal is to devise a set of practices that allow an AS concerned about routing security risk (in particular the risk of BGP hijack) to take action that minimizes this risk. In other words, they will shift interconnection patterns to exploit the benefit of the zone.

For many small to mid-size networks, connecting to an Internet exchange (IX) is an efficient way to establish many peering connections. For our hypothetical zone with 900 members, of the 57 288 customers of the zone, we identify 15 337 that attach to at least one IX. Some IXes may take steps to reduce the risk of hijacks among their members, such as requiring that their members document the ASs that they will legitimately announce. Some peers may take steps to verify their own customers, and the practical risk of using routes

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

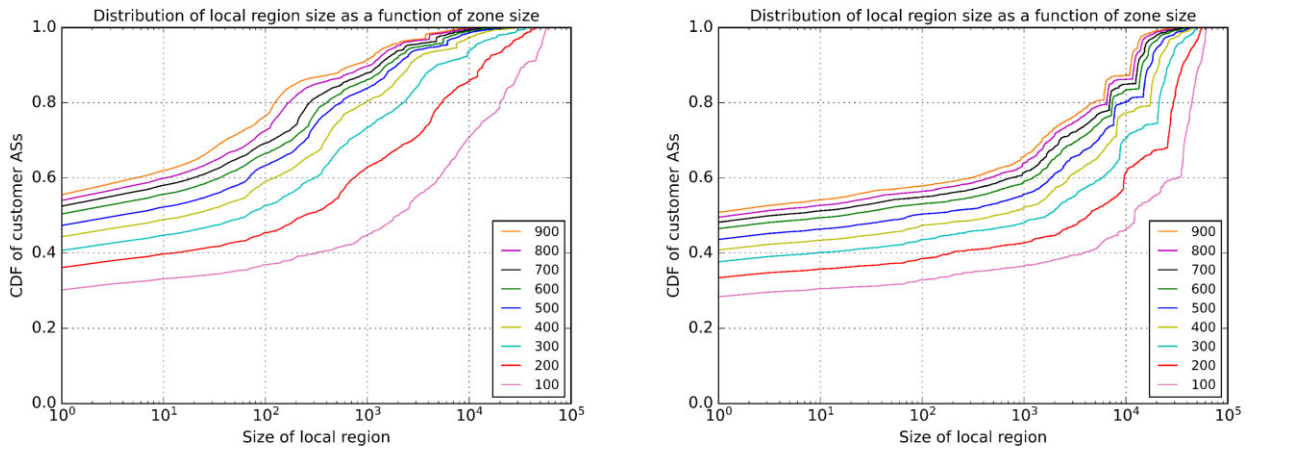

(a) Plot based on data from ASrank

(b) Plot based on data from ASrank plus an assumed peering relationship between any two ASes at an IX

Figure 6. Local region sizes for customers of a hypothetical VIPzone, for various zone sizes. Between 30% and almost 60% of the customer ASes have a region size near 1.

from such a peer would be minimal. (Internet2 exemplifies such a region; they track the full customer cone of their members, and use prefix filters to prevent incorrect announcements. Using routes from a region of this sort is practically risk-free.) It may not be practical for every AS that connects to an exchange to assess the pragmatic level of risk at that exchange, but if one actor can make that assessment on behalf of the exchange, all the members can take that information into account in deciding what action to take, e.g. whether to peer with all of the other exchange members, or peer selectively with those peers that offer significant traffic volumes.

A further uncertainty in Fig. 6 derives from the common use of prefix filters on peering links, precisely to protect themselves from harm due to erroneous or malicious BGP announcements. That practice would reduce the effective size of the local region from which hijacks can come. Many ASes consider such filtering good routing hygiene today [88].

But we emphasize that the power of a trust zone approach is that each AS gets to make its own risk assessment, and act accordingly. ASes with small or no local region would not have to take these steps. Larger ASes are more likely to have the operational capacity to protect their local region, e.g. implement prefix filters. If a VIPzone existed, we would expect ASes to take actions to reduce the residual risks from their local regions.

Protection for ASes not attached to the zone

In Fig. 5, AS B shares the local region of A, but is not directly connected to the zone. What protection does B receive from hijacks? With respect to owner risk, B can prevent simple hijacks based on an invalid origin by registering ROAs, but it gets no protection from path hijacks. With respect to misdirection risk, it is in the same situation as A: no hijacks will come into B's region from the zone, but a hijack in B's local region can still cause misdirection harm. Many smaller ASes offer low-value, limited-interest services, and their owner risk of a hijack is minimal. If the AS does consider the owner risk to be substantial, they can and should obtain transit from a member of the zone.

Auditing requirements

Our proposal for a VIPzone does not use real-time detection of suspicious announcements. Real-time prevention requires adding code to the BGP processing path in routers or route computation servers. This approach would potentially lead to a more brittle scheme.

Instead, the VIPzone uses a trust-but-verify approach: checking conformance of members with its requirements, detection and documenting of failures, and suspension or ejection of non-compliant members. This requirement means that members must have the will, and the institution, to undertake conformance auditing. Independent third parties can check conformance off-path, by looking at public BGP announcements. In support of this auditing, every VIPzone member would be required to provide a BGP view to a route collector. The audit process does not use the member's view to audit their behavior (the member could lie) but rather uses the views provided by the member's neighbors that are also members and thus provide views of their own. Using the neighbor views allows confirmation that the member correctly propagated verified routes with the VERIFIED tag, and did not use the VERIFIED tag on routes that other members had not tagged as VERIFIED.

This approach is similar in spirit to how the CA/Browser forum verifies the correct behavior of certificate authorities. Its goal is not to detect and block every issuance of a false certificate in real time, but rather to identify CAs that are shown to be untrustworthy so that providers of browsers can choose to remove them from their list of trusted root CAs. The idea is to enforce proper behavior by making the consequence of misbehavior a substantial penalty. In that context, the CA/Browser community has shown a willingness to take action against providers that do not conform. For the VIPzone to provide protection in practice, the routing community must have the same will. We argue that an industry-led body, analogous to the CA/Browser forum but with a stronger authority, should decide on necessary actions if a VIPzone member does not conform to the required practices. But note that the penalty in this case is not being disconnected from the Internet, but just losing the right to initiate VERIFIED announcements.

Independent of the exact specification of the practices that define a zone, it must be possible to tell by inspection if an announcement is not conformant. The three tests for VIPzone member conformance are:

(i) Rule 1: If an announcement (observed anywhere in the VIPzone) has more than one AS number in the path before it enters the VIPzone, and is marked VERIFIED, the member that in-

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

troduced the announcement into the core is non-conformant. Our trust model assumes that verification and checking of announcements occurs at specific locations: the ASes at the edge of the VIPzone that have customers not in the VIPzone. This requirement makes it possible to identify members that do not implement the required practices.

(ii) Rule 2: If an announcement has an invalid origin, as determined by a ROA, independent of path length, the VIPzone member that introduced the announcement is non-conformant.

(iii) Rule 3: ASes in the VIPzone must forward the VERIFIED community value from other VIPzone members.

Another advantage of off-line conformance checking is that it could allow an AS to register its intent to violate Rule 1 in a specific case and announce a route that is non-conformant (e.g. to deal with a specific customer requirement), and accept responsibility for ensuring that it is benign. Allowing benign exceptions, including marking them VERIFIED, enables more nuanced balance of security and availability priorities.

Economic cost of checking for conformance

The conformance checking requirements imply non-trivial costs. The data collection and curation infrastructure would require staffing to maintain, whether operated by an independent private-sector group such as RouteViews, or some more formally chartered institution or agency. Then, technically capable organizations must perform the auditing and provide the information necessary to judge untrustworthy behavior.

Comparison to other proposed solutions

We compare our proposal to two leading alternative proposals to advance the collective state of routing security, in particular to prevent path hijacks: BGPsec and ASPA. But we preface this comparison with a comment on the tension between our VIPzone approach and the philosophy of *zero trust architectures*. Zero trust is usually proposed in a context where each machine or subsystem performs its own verification to protect itself, and the incentives are directly aligned [89]. The collective action aspect of routing security, where it is not feasible to verify implementation by other parties, is at odds with this assumption.

The VIPzone approach allocates responsibility to specific points in the zone (the perimeter), and ascertains whether zone members are implementing the required operational practices.

Autonomous System Provider Authorization

ASPA [74] is a mechanism that lets a customer AS register a list of providers that the customer uses. The AS registers an Autonomous System Provider Authorization or ASPA record in the same system that stores ROAs—the RPKI administered by the five RIRs. The ASPA data is globally visible, so any AS receiving a BGP announcement can look at the sequence of ASes in the path, and check to see if there is an ASPA that covers any adjacent pair of ASes in the path. If there is, and the announcement is inconsistent with the ASPA, the AS receiving the announcement can drop it [74]. ASPA can be used to limit both route leaks and, to some degree, against path hijacks, assuming the appropriate ASes deploy ASPA in the correct places. The ASPA specification describes several deployment scenarios.

ASPA's design differs in several ways from our proposal.

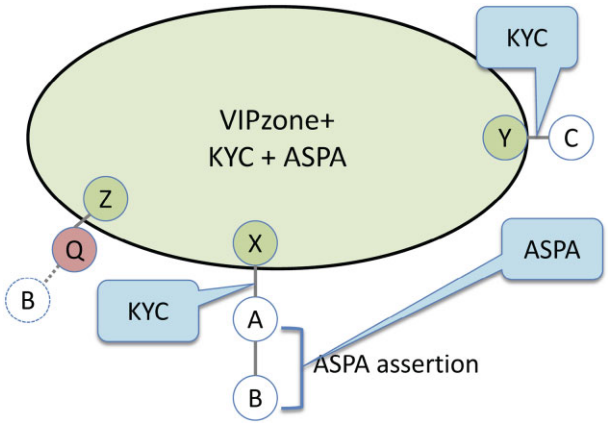

Figure 7. ASPA-extended VIPzone. Zone member X can use ASPA to verify an announcement with two ASes in the path outside the zone. If B registers an ASPA recording A as a provider, and X has done KYC on A, then X can mark the route as VERIFIED, extending the scope of the VIPzone.

(i) The VIPzone design tries to minimize the effort required of small ASes to get protection. It requires only that the small AS connect to a transit provider that is in the zone and (ideally) register its ROAs. ASPA requires that the small AS register an ASPA describing its providers. While the mechanics of registration need not be complex, this registration becomes one more data record that the operators of the AS must keep track of, and remember to change if they change providers.

(ii) The VIPzone design does not require new mechanisms in the routers (or route computation servers). The actions required of a VIPzone member (see our discussion in Verified IP Zone) include new operational practices and use of a new BGP community value. ASPA checking requires a new processing check, which includes downloading the relevant ASPA data and inspecting the announcement for validity. This dependency also implies the need for the RPKI to store and manage new (ASPA) records.

(iii) The VIPzone design assigns clear responsibilities: an AS at the edge of the zone has specific requirements to check announcements received from its customers, including a KYC check. This perimeter allows clear description of protection and residual harm. The current ASPA draft [74] describes use cases without assigning responsibilities to specific ASes. Thus it is not clear which ASes should do ASPA checking, which ASes would have the motivation to register ASPA records, and (thus) what protection ASPA will achieve. For example, if an AS has listed a provider in an ASPA record, and that provider has such poor business/operational practices that it cannot identify an imposter posing as their legitimate customer, an ASPA record alone cannot prevent the resulting harm. Assignment of responsibility, as in VIPzone, allows the possibility of conformance checking.

ISPs could use ASPA to extend the range of VIPzone protection to customers of zone customers. This extension would allow an AS at the edge of a zone to mark as VERIFIED those announcements with two or fewer ASes in the path, as opposed only one. To illustrate, in Fig. 7, Y uses VIPzone practices to verify announcements originated by C, as does X to verify announcements from A. But X has the option of using an ASPA registered by B to confirm that A is a valid provider of B. X can tag as VERIFIED the announcement that includes both A and B in the path only if there is an ASPA registered by B. Otherwise,

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

X must not mark the announcement, but can forward it into the zone unmarked.

The other case in Fig. 7 is that Q is malicious, and wants to hijack a prefix belonging to B. If B has registered a ROA for the prefix, then Q cannot validly announce B's prefix. It would have to pretend to be B. If B has registered an ASPA saying that its provider is A, then this ASPA would allow Z to conclude that Q's announcement is invalid. The attacker Q could add AS A to the path to make a valid path, but then the announcement would have three ASes in the path outside the zone (A, B, and Q), and (in the VIPzone we propose) Z must not mark a path VERIFIED if it has more than two ASes in the announcement.

BGPsec

Like ASPA, BGPsec is attempting to achieve a zero-trust approach. To the extent that every router that forwards the announcement adds its own cryptographic signature, any router along the path can verify that the series of signatures to that point are valid. This function also means that BGPsec, if pervasively and correctly deployed, provides the technical means to address the KYC requirement that VIPZone and ASPA cannot do in-protocool. That is, BGPsec prevents the social engineering impersonation attack, assuming the imposter does not have the necessary keys to sign their announcements. However, the requirement for comprehensive deployment dramatically reduces the incentive for ISPs to undertake the cost and complexity of BGPsec deployment, and a lengthy trajectory of partial deployment implies inconsistent and unpredictable implementation of the required checking. We expect that governments will not have the patience to wait for deployment of a global solution to route hijacks.

Conclusion

There is currently no consensus as to the next step to secure BGP beyond the simplest type of hijacks. As of 2024, BGPsec has no production deployment and arouses significant controversy over the operational feasibility of its key management aspects. For all proposed solutions to prevent path hijacks, incentives are misaligned. We have proposed a path forward that creates incentives for ASes (both customer and provider) to participate, protects ASes against *path hijacks* and *origin hijacks* with no effort or investment needed by small ASes, and avoids the need for new mechanism in routers.

One insight that shapes our proposal is that if there is a coherent topological region of the Internet, and with practices limiting malicious BGP routes entering that region, then the operational practices can provide much stronger protection against abuse for those who join, and thus incentive to participate. The result is a virtuous circle, where customers benefit from choosing ISPs committed to the practices, and ISPs (thus) benefit from committing to the practices. A coherent core of ISPs has already emerged organically in the ecosystem, which can be leveraged to create a *zone of trust*, a region that protects not only all networks in the region, but all directly attached customers.

A few concerns with VIPzone bear further consideration. First, will it concentrate power in a few trusted networks, those with the authority to verify routes? We believe the VIPzone requirement for transparency, accountability, and independent auditing, provides a counterpoint to potential abuses of power.

Second, will trust zones fragment the Internet? Some Internet fragmentation has already occurred, and trust zones provide a way to bridge some of these fragments using a trust-but-verify framework, like treaties in other global domains. We acknowledge that multiple trust zones may emerge, including on national boundaries. But note that for a VIPzone to be effective, both (1) ASes that produce important services and (2) ASes that consume those important services, must be attached to that zone. As an extreme example, if each country wants its own trust zone, networks with global customer bases would have to replicate their point of attachments in all trust zones where they serve customers. We imagine trust zones to evolve instead more like global trading zones.

Third, achieving the VIPzone protection requires auditing and enforcing conformance with the practices. The institutional framework required for the necessary data collection already exists in multiple places, e.g. RIPE and RouteViews. But it is still more expensive (and therefore less incentive-compatible) than doing nothing in the current unregulated environment.

Fourth, ISPs have to trade off some autonomy in exchange for routing security. ISPs are required to prefer VERIFIED routes over customer routes, and ISPs would hand some control over to a non-ISP third party (the auditor) similar to the CA/Browser Forum today. But unlike other proposed approaches to routing security, transit ISPs can claim to offer their customers a securely routed service by participating.

Our proposal responds to a long-standing need for some medium-term path forward on protection against path hijacks. We believe it is a direction worth debate and analysis in the context of possible regulatory measures. We recognize that ISPs, like most private sector actors, prefer lack of regulation and work to avoid it as long as possible. But the EU has made it clear they will regulate to safeguard their citizens despite private sector objections [90–92]. We offer this path forward as an interdomain routing approach for which the private sector could drive a self-regulatory framework that achieves the accountability regulators are now seeking in digital domains.

Acknowledgments

The authors thank the anonymous reviewers for their valuable suggestions. The views and conclusions included here are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the US National Science Foundation.

Author contributions

David Clark (Conceptualization, Investigation, Validation, Visualization, Writing—original draft, Writing—review & editing), Cecilia Testart (Conceptualization, Investigation, Writing—original draft, Writing—review & editing), Matthew Luckie (Conceptualization, Data curation, Investigation, Validation, Visualization, Writing—original draft, Writing—review & editing), and kc claffy (Conceptualization, Funding acquisition, Writing—original draft, Writing—review & editing)

Conflict of interest: None declared.

Funding

This work is supported in part by funds from the National Science Foundation (NSF grant OAC-2131987 and NSF grant CNS-2120399).

Appendix 1. Full specification of required actions for members of the VIPzone

We summarize the required operational practices of VIPZone members in the body of the paper; here we repeat the summary, provide additional details, and diagram specific scenarios.

First, VIPzone members that can participate in these enhanced practices must be part of a connected region.

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

Second, if a VIPzone member receives a BGP announcement from a neighbor that is not in the zone, and the announcement is for a prefix that the neighbor *originates* and the member can verify as legitimate, then the member will tag the route with a new BGP community value [13], which we call *VERIFIED*. (Some other BGP mechanism with equivalent properties could also be used.)

Third, VIPzone members must propagate this community value as they forward announcements to other ASes. This allows neighbors to establish the authenticity of the route, regardless of the distance they are from the origin.

Fourth, inside the zone, any AS receiving multiple announcements for the same prefix must prefer one marked VERIFIED. By this rule, no member will prefer a path hijack announcement over a legitimate announcement from customers directly attached to the zone, since legitimate announcements will be marked VERIFIED.

The operational practices that a VIPzone member must configure their routers to follow are as given below:

(1) **Prevent false VERIFIED routes:** If the member receives an announcement from a non-member AS, then it MUST remove the VERIFIED community if present. This is to prevent an attacker from injecting a hijacked route that other VIPzone members prefer.
(2) **Drop RPKI-invalid routes:** If the member receives an announcement where the origin is RPKI-invalid, the member MUST drop the announcement. This step prevents origin hijacks.
(3) **Prevent propagation of forged routes:** If the member receives an announcement where the AS used by the neighbor is not consistent with the AS numbers legitimate for the neighbor, the member MUST drop the announcement. This step implements a Know Your Customer (KYC) requirement, to prevent malicious routes from entering the VIPzone.
(4) **Forward VERIFIED routes:** If the member receives from another member an announcement with a VERIFIED community tag set, it MUST retain that tag when forwarding the route to other members. Further, the member MUST retain the VERIFIED tag when it provides the route to non-member neighbors. Customers of zone members do not need to understand or act on the VERIFIED marking; this zone rule allows them the option to distinguish which routes have been VERIFIED on entry to the zone, and thus are not path hijacks.
(5) **Verify routes with one AS in the path from non-member customers:** If the member receives from a non-member customer an announcement with one AS in the path, the member MUST drop the announcement if the route contains a prefix that the customer has no authority to announce (it is not RPKI-valid, or is not from a list of prefixes that the member has previously established as allowed from their customer). If the prefix is RPKI-valid, is registered by the owner in an authenticated IRR, or from a list of allowed prefixes, the zone member AS MUST add a VERIFIED-community to the route before propagating it, so that other members know that the route is valid.
(6) **Forward unverified routes without the VERIFIED tag:** If the zone member has not established that the announcement is valid (because it has not yet obtained the list of allowed prefixes or because the AS path in the route contains more than one unique ASN and so cannot be verified), the member can announce the route to its neighbors but MUST NOT add a VERIFIED community to the route, so that other members know not to trust the validity of the route. To preserve Internet connectivity, zone members must forward unverified routes according to normal routing policies.
(7) **Export routes to a route collector for auditing:** To allow for auditing the behavior of trust zone members, members must export their routes to a route collector.

Appendix 2. Hijack scenarios in a local region

We elaborate on some implications of a local region.

Multihoming transit scenario

VIPzone members, and non-members exclusively connected to VIPzone transit providers, will receive an authentic route from the VIPzone if one is available. The hijack risk for a VIPzone customer increases if the customer accepts routes from a non-member in its local region. Figure B1 illustrates the scenario of the residual risk in a local region from a transit provider that is not in the zone. Here AS J connects to two transit providers (M_5 and X), of which only M_5 is a zone member. AS X and AS Z are in the local region of J, since they can originate BGP announcements that arrive at J without passing through the zone. If X or Z sends a bogus announcement for a prefix (in this example, B) to J, J might decide to prefer it over a valid (VERIFIED) route from M_5. This could happen only if X or Z are malicious—given the local region of J, there are no other ASes in a position to launch a hijack.

If J did not use X as a transit provider, or preferred the VERIFIED route from M_5, it would prevent this hijack.

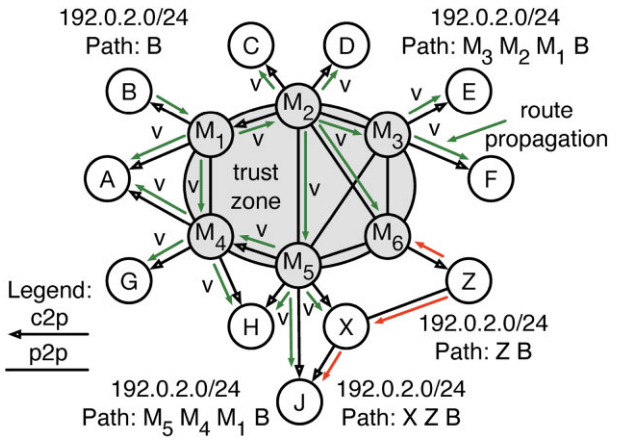

Figure B1. Multhoming scenario. The risk of a hijack of J's traffic, in this case destined to prefix B, is limited to BGP announcements coming from X or Z, but that assumes that J chooses that route rather than a VERIFIED route to B from its zone transit provider M_5.

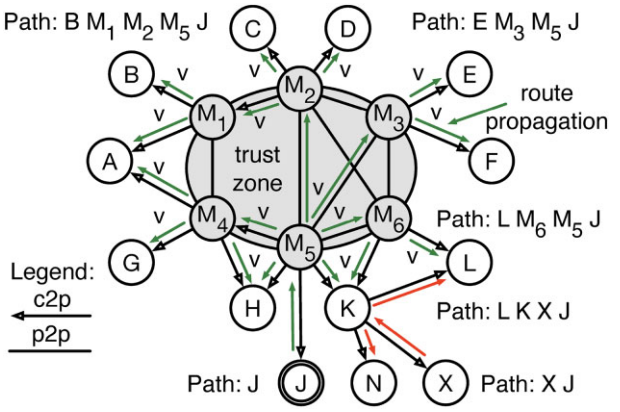

Figure B2. An illustration of how protection depends on how customers connect to the Routing Trust Zone. Customer L receives two routes for J's prefix, a VERIFIED route via M_6 and an (unverified) hijacked route from X via K. If L does not prefer the VERIFIED route via M_6 (L M_6 M_5 J), L may select the hijacked route (L K X J) because it has the same AS path length.

Use of VERIFIED outside the zone

While the VIPzone practices are not required or expected for non-members, a non-member may choose to configure their routers to remove VERIFIED tags from non-member neighbors, and then prefer routes received from their neighbors who are VIPzone members that are tagged as VERIFIED, to avoid using a malicious hijacked path towards a destination (misdirection harm). VIPzone members must retain VERIFIED tags so that non-members could select these routes.

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

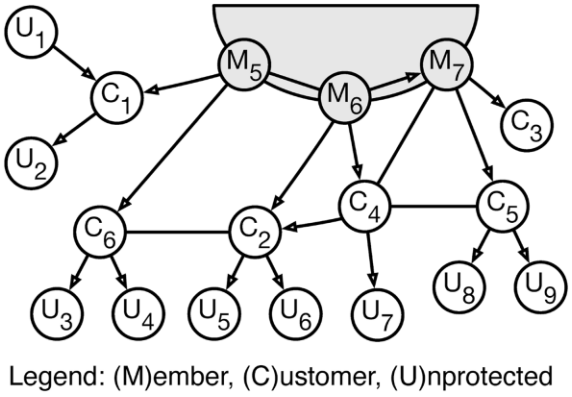

Figure B3. Peering across the VIPzone perimeter. C_2 has two transit providers, only one of which (M_6) is in the VIPzone. M_6 will announce into the VIPzone a verified path to C_2. C_4 peers with M_7, which is in the VIPzone. C_4 will announce to M_7 a customer route to C_2. M_7 will prefer the VERIFIED announcement, and will send traffic to C_4 through its provider M_6 in the zone, not over the peering link to C_4.

In Fig. B2, L could have chosen a VERIFIED route via M_6 if L preferred routes received from VIPzone members that are VERIFIED over unverified routes. L has an incentive to do so, as otherwise when it receives a path hijack by X of equal length (e.g. L K X J from K) to the authentic route via its VIPzone provider, it may select the hijacked route and suffer associated harms. Note that L has a higher risk of accepting a hijacked route from a peer or customer AS, as those routes ordinarily have a higher preference than a provider route.

Peering interconnection scenarios

In most cases, the analysis for a peering connection is similar to transit connections.

Peering with IXP route servers

An IXP-operated route server centralizes peering routes from IXP members and makes these routes available to other IXP members. If the IXP is a member of the VIPzone and has configured the route server to verify routes received from IXP members, then the route server can mark routes as VERIFIED, and VIPzone members can propagate the VERIFIED route. Otherwise, routes received from a route server are unverified. (At least one IXP (INEX) has been performing ROV filtering on its route server since February 2019 [93].)

Peering of zone customers outside zone

If two ASes not in the VIPzone but directly connected to VIPzone providers peer with each other, they may receive announcements of routes to each other via the VIPzone that are marked VERIFIED, and announcements over the peering connection that are not VERIFIED. Because ASes not in the VIPzone are not expected to use that community value to assign a preference to an announcement, their routing policy would be the same as today. This case illustrates a local region. (Note ASes outside the zone may choose to use this VERIFIED value to prefer routes. but they had better know what they are doing, because it may cause unexpected results, e.g. use of paths via a provider rather than a peer.)

Peering across the VIPzone perimeter

Peering across the VIPzone perimeter has a straightforward scenario and a complicated scenario. Imagine that VIPzone member M_7 in Figure B3 peers with non-VIPzone C_4. In the straightforward case, M_7 will apply the same VIPzone rules to peer C_4 as it does for customers C_3 and C_5, i.e. forward or drop announcements and mark as VERIFIED announcements that peer C_4 legitimately originates. Customers of that peer C_4 would not have their routes VERIFIED. Typical routing policy is that the AS in the zone would only use these announcements from peer C_4 for itself and its customers—it would not forward them on to other peers or providers.

The complicated peering scenario arises when a customer of that non-zone member also obtains transit service from an AS in the VIPzone. Figure B3 shows C_2 with two transit providers, only one of which is in the zone. The transit provider not in the zone (C_4) also peers with an AS in the zone (M_7). In this case, M_7 will receive a VERIFIED announcement to C_2 via M_6, which per the VIPzone rules it must prefer over the route via the peering link from C_4, so M_7 will not benefit from the peering link (through C_4) for traffic to C_2, even if it would normally prefer that peering link.

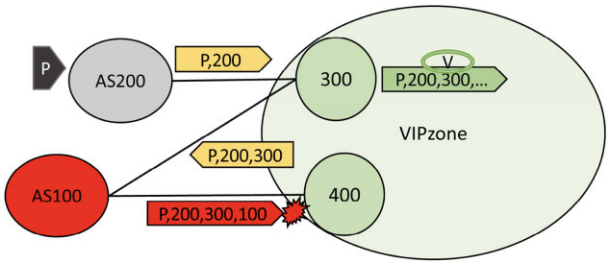

Figure C1. AS 100 legitimately receives from its provider AS300 in the zone a route to prefix *P* in AS200. If AS100 leaks this route to AS400, the route [which includes multiple AS hops (200, 300, 100)] will not be marked VERIFIED, and since there is a verified announcement for the prefix *P*, S400 will not prefer the leaked route. The leak has no effect.

Appendix 3. Route leaks

As mentioned in our discussion of the VIPzone, a route leak is an event in which an AS inappropriately (i.e. violating routing policy) forwards a route it legitimately received. The consequence is often that large flows of traffic reach this AS, which is not provisioned to carry them. A classic route leak occurs when a multi-homed AS that takes the routes it receives from one of its transit providers and inadvertently propagates these routes to its other transit provider (Fig. C1).

In addition to preventing path hijacks of ASes directly attached to the zone, the VIPzone prevents leaks of announcements of prefixes belonging to those ASes (Fig. C1). AS 100 might incorrectly announce (leak) the path to AS 200 that it receives from one transit provider (AS 300) to its other transit provider (AS 400). Since a VERIFIED path to AS 200 exists in the zone, AS 400 should not propagate its unverified route. If AS 400 did propagate its unverified route, ASes in the VIPzone would never prefer that route, so customers directly attached to the VIPzone would not receive that route, and traffic to AS 200 would never flow from the zone to AS 400.

The VIPzone we have constructed protects against route leaks by ASes *not* in the VIPzone. If the leak occurs within the zone, the announcement from AS 300 to AS 100 would be VERIFIED, and when AS 100 forwards (leaks) this announcement to AS 400, AS 400 must remove the VERIFIED marking first if it propagates the announcement to the zone.

Such potential harms from accidental misconfiguration suggest an important insight about VIPzone deployment. A natural but unnecessary—even counterproductive—objective is to maximize the number of ASes in the VIPzone. Smaller ASes (certainly stub ASes) will get the benefit of VIPzone from being a *customer* of a VIPzone member. Actually joining will require that the joining AS correctly implement a range of operational practices, which for smaller ASes with less sophisticated staff may be difficult. Getting these practices wrong may result in malformed announcements in the zone, which will lead to the revocation of their VIPzone status. We consider it preferable that only operators with sufficient technical abilities attempt to join the VIPzone. Other requirements (such as maintaining correct contact information, registering their own prefixes in a public database, and implementing anti-spoofing filters) make sense for an AS of any size, and a MANRS-like initiative may want to define two tiers of ISP membership to accommodate different likely capabilities.

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

References

1. Apostolaki M, Zohar A, Vanbever L. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In: *2017 IEEE Symposium on Security and Privacy (SP)*. San Jose, California, USA: Institute of Electrical and Electronics Engineers (IEEE), 2017.
2. Goodin D. Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency. 2018. https://arstechnica.com/information-technology/2018/04/suspicious-event-hijacks-amazon-traffic-for-2-hours-steals-cryptocurrency/ (1 September 2023, date last accessed).
3. Dan Goodin. How 3 hours of inaction from Amazon cost cryptocurrency holders $235,000. 2022. https://arstechnica.com/information-technology/2022/09/how-3-hours-of-inaction-from-amazon-cost-cryptocurrency-holders-235000/ (1 September 2023, date last accessed).
4. RIPE Network Coordination Centre. YouTube Hijacking: A RIPE NCC RIS case study. 2008. https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study (1 September 2023, date last accessed).
5. Paganini P. BGP hijacking - Traffic for Google, Apple, Facebook, Microsoft and other tech giants routed through Russia. 2017. https://securityaffairs.co/wordpress/66838/hacking/bgp-hijacking-russia.html (1 September 2023, date last accessed).
6. Ramachandran A, Feamster N. Understanding the network-level behavior of spammers. In: *ACM SIGCOMM Computer Communication Review*. Vol. **36**. New York, New York, USA: ACM, 2006.
7. Vervier PA, Thonnard O, Dacier M. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In: *Proceedings 2015 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2015.
8. White Ops and Google. The Hunt for 3ve: Taking down a major ad fraud operation through industry collaboration. Mountain View, California, USA: Google, 2018.
9. Snijders J. The BGPsec plan. 2022. https://www.bgpsec.net/history.html (1 September 2023, date last accessed).
10. U. S. Federal Communications Commission. NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing. 2022. https://www.fcc.gov/ecfs/document/1022806680214/1 (1 September 2023, date last accessed).
11. U. S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency (CISA). NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing. 2022. https://www.fcc.gov/ecfs/document/1022806680214/1 (1 September 2023, date last accessed).
12. U. S. Department of Justice National Security Division (Matthew G. Olsen) and U. S. Department of Defense Acquisition and Sustainment (William A. Laplante). Public Comment to NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing. 2022. https://www.fcc.gov/ecfs/document/1091496862125/1 (1 September 2023, date last accessed).
13. Chandra R, Traina P, Li T. BGP Communities Attribute. IETF RFC 1997, 1996. https://www.rfc-editor.org/rfc/rfc1997.html (1 September 2023, date last accessed).
14. Donnet B, Bonaventure O. On BGP Communities. *ACM CCR* 2008; **38**: 55–9.
15. King T, Dietzel C, Snijders J. *et al.* RFC 7999: BLACKHOLE Community. 2016, https://datatracker.ietf.org/doc/html/rfc7999 (1 September 2023, date last accessed).
16. CISCO. Remotely Triggered Black Hole Filtering - Destination Based and Source Based. 2005, Cisco White Paper, http://www.cisco.com/c/dam/en_us/about/security/intelligence/blackhole.pdf (1 September 2023, date last accessed).
17. Gao L. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans Networking* 2001;**9**:733–45.
18. Rosen E. Exterior Gateway Protocol (EGP), RFC 827. 1982. https://www.rfc-editor.org/info/rfc827 (1 September 2023, date last accessed).
19. Merit. IRR - Internet Routing Registry. http://www.irr.net (1 September 2023, date last accessed).
20. Ronald F. Guilmette. Cogent and FDCServers: Knowingly aiding and abetting fraud and theft? 2019. https://mailman.nanog.org/pipermail/nanog/2019-September/102963.html (1 September 2023, date last accessed).
21. Ostap Efremov. 196.52.0.0/14 revoked, cleanup efforts needed. 2021. RIPE NCC Anti-Abuse Working Group. https://www.ripe.net/participate/mail/forum/anti-abuse-wg/PENBT0dHenF6bUcwd1VOYnNZbz1oamQ9K1JjTmFvc09XR0xOMGpxV0JnVEpteFBocFItQUBtYWlsLmdtYWlsLmNvbT4= (1 September 2023, date last accessed).
22. Toonk A. Using BGP data to find Spammers. 2014. New York, New York, USA: Association for Computing Machinery (ACM). https://bgpmon.net/using-bgp-data-to-find-spammers/.
23. Du B, Akiwate G, Krenc T. *et al.* IRR Hygiene in the RPKI Era. In: *PAM*. New York, New York, USA: Association for Computing Machinery (ACM), 2022, 321–37.
24. Du B, Izhikevich K, Rao S. *et al.* IRRegularities in the Internet Routing Registry. In: *Proceedings of the 2023 ACM on Internet Measurement Conference. IMC '23*, New York, NY: Association for Computing Machinery, 2023, 104–10.
25. Oliver L, Akiwate G, Luckie M. *et al.* Stop, DROP, and ROA: Effectiveness of Defenses through the Lens of DROP. In: *Internet Measurement Conference*. Nice, France: Association for Computing Machinery (ACM), 2022.
26. Heffernan A. RFC 2385: Protection of BGP Sessions via the TCP MD5 Signature Option. 1998. https://tools.ietf.org/html/rfc2385 (1 September 2023, date last accessed).
27. Touch J, Mankin A, Bonica RP. RFC 5925: The TCP authentication option. 2010. https://tools.ietf.org/html/rfc5925 (1 September 2023, date last accessed).
28. Smith BR, Garcia-Luna-Aceves JJ. Securing the Border Gateway Routing Protocol. In: *Global Telecommunications Conference*. London, United Kingdom: Institute of Electrical and Electronics Engineers (IEEE), 1996.
29. Kent S, Lynn C, Seo K. Secure Border Gateway Protocol (S-BGP). *IEEE J Sel Areas Commun* 2000;**18**:582–92. https://doi.org/10.1109/49.839934.
30. Gouda MG, Elnozahy EN, Huang CT. *et al.* Hop integrity in computer networks. *IEEE/ACM T Networking* 2002;**10**:308–19. https://doi.org/10.1109/TNET.2002.1012363.
31. Zhao X, Pei D, Wang L. *et al.* Detection of invalid routing announcement in the Internet. In: *Proceedings International Conference on Dependable Systems and Networks*. London, United Kingdom: Institute of Electrical and Electronics Engineers (IEEE), 2002.
32. White R. Securing BGP through secure origin BGP. *Internet Protoc J* 2003;**6**:47–53.
33. Goodell G, Aiello W, Griffin T. *et al.* Working around BGP: an incremental approach to improving security and accuracy in interdomain routing. In: *ISOC Symposium on Network and Distributed Systems Security*. San Diego, California, USA: The Internet Society, 2003.
34. Hu YC, Perrig A, Sirbu M. SPV: secure path vector routing for securing BGP. *ACM SIGCOMM Comput Commun Rev* 2004;**34**:179–92.
35. Subramanian L, Roth V, Stoica I. *et al.* Listen and whisper: security mechanisms for BGP. In: *Symposium Networked System Design and Implementation*. San Francisco, California, USA: USENIX: The Advanced Computing Systems Association, 2004.
36. Wan T, Kranakis E, van Oorschot PC. Pretty secure BGP, psBGP. In: *Proceedings of the 2005 ISOC Symposium on Network and Distributed Systems Security*. The Internet Society: San Diego, California, USA, 2005.
37. Nordstrom O, Dovrolis C. Beware of BGP attacks. *ACM CCR* 2004;**34**:1–8.
38. Karlin J, Forrest S, Rexford J. Pretty good BGP: improving BGP by cautiously adopting routes. In: *IEEE International Conference on Network Protocols*. Santa Barbara, California, USA: Institute of Electrical and Electronics Engineers (IEEE), 2006.
39. Reynolds P, Kennedy O, Sirer EG. *et al.* Using external security monitors to secure BGP. In: *Computing and Information Science, Technical Report TR2006-2065*. Cornell University, 2006.

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

40. Qiu J, Gao L. Hi-BGP: a lightweight Hijack-proof Inter-domain Routing Protocol. 2006. https://www.researchgate.net/publication/228871049_Hi-bgp_A_lightweight_hijack-proof_inter-domain_routing_protocol (1 September 2023, date last accessed).
41. Israr J, Guennoun M, Mouftah HT. Credible BGP–Extensions to BGP for Secure Networking. In: *Fourth International Conference on Systems and Networks Communications*. Porto, Portugal: Institute of Eletrical and Electronics Engineers (IEEE), 2009.
42. Lepinski M, Sriram K. RFC 8205: BGPsec Protocol Specification. 2017. https://tools.ietf.org/html/rfc8205 (1 September 2023, date last accessed).
43. Nicholes MO, Mukherjee B. A survey of security techniques for the Border Gateway Protocol (BGP). *IEEE Commun Surv Tutor* 2009;**11**:52–65. https://doi.org/10.1109/SURV.2009.090105.
44. Butler K, Farley TR, McDaniel P. *et al.* A Survey of BGP Security Issues and Solutions. *Proc IEEE* 2010;**98**:100–22.
45. Siddiqui MS, Montero D, Serral-Gracia R. *et al.* A survey on the recent efforts of the Internet Standardization Body for securing inter-domain routing. *Comput Networks* 2015;**80**:1–26. https://doi.org/10.1016/j.comnet.2015.01.017.
46. Mitseva A, Panchenko A, Engel T. The state of affairs in BGP security: a survey of attacks and defenses. *Comput Commun* 2018;**124**:45–60. https://doi.org/10.1016/j.comcom.2018.04.013.
47. Testart C. Reviewing a historical Internet vulnerability: why isn't BGP more secure and what can we do about it?. In: *Telecommunications Policy Research Conference*. Washington D.C., USA: SSRN, 2018.
48. Gill P, Schapira M, Goldberg S. Let the market drive deployment: a strategy for transitioning to BGP security. *ACM SIGCOMM Comput Commun Review* 2011;**41**:14–25. https://doi.org/10.1145/2043164.2018439.
49. Lychev R, Goldberg S, Schapira M. BGP security in partial deployment: is the juice worth the squeeze?. In: *SIGCOMM*. Hong Kong, China: Association for Computing Machinery (ACM), 2013, 171–82.
50. Goldberg S. Why is it taking so long to secure Internet routing?. *Comm ACM* 2014;**57**:56–63.
51. Cohen A, Gilad Y, Herzberg A. *et al.* Jumpstarting BGP security with path-end validation. In: *ACM SIGCOMM*. Florianopolis, Brazil: Association for Computing Machinery (ACM), 2016.
52. Zhang X, Hsiao HC, Hasker G. *et al.* SCION: scalability, control, and isolation on next-generation networks. In: *2011 IEEE Symposium on Security and Privacy*. Oakland, California, USA: Institute of Electrical and Electronics Engineers (IEEE), 2011, 212–27.
53. Birge-Lee H, Wanner J, Cimaszewski GH. *et al.* Creating a Secure Underlay for the Internet. In: *USENIX Security Symposium*. Boston, Massachusetts, USA: USENIX: The Advanced Computing Systems Association, 2022.
54. Swisscom Internet Exchange. SCION Peering Partcipants 2024. https://www.swissix.ch/services/scion-peering-mesh/scion-peering-participants/ (1 September 2023, date last accessed).
55. Scudder J, Bush R, Mohapatra P. *et al.* RFC 6811: BGP Prefix Origin Validation. 2013. https://tools.ietf.org/html/rfc6811 (1 September 2023, date last accessed).
56. Huston G, Michaelson G. RFC 6483: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs). 2012. https://tools.ietf.org/html/rfc6483 (1 September 2023, date last accessed).
57. Gilad Y, Goldberg S, Sriram K. *et al.* The Use of maxLength in the Resource Public Key Infrastructure (RPKI) (RFC9319). 2022. https://www.rfc-editor.org/info/rfc9319 (1 September 2023, date last accessed).
58. KPN. AS286 Routing Policy. https://as286.net/AS286-routing-policy.html (1 September 2023, date last accessed).
59. Borkenhagen J. AT&T/as7018 now drops invalid prefixes from peers. 2019. https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html (1 September 2023, date last accessed).
60. Telia. Telia Carrier Takes Major Step to Improve the Integrity of the Internet Core. 2019. https://www.businesswire.com/news/home/20190915005013/en/ (1 September 2023, date last accessed).
61. Jason Livingood. 2021. https://corporate.comcast.com/stories/improved-bgp-routing-security-adds-another-layer-of-protection-to-network (1 September 2023, date last accessed).
62. NIST. NIST RPKI Monitor 2.0: Methodology and User's Guide. 2022. https://rpki-monitor.antd.nist.gov/Methodology#ROV_Donut (1 September 2023, date last accessed).
63. National Institute for Stantards and Technology. RPKI Deployment Monitor. 2022. https://rpki-monitor.antd.nist.gov/ (1 September 2023, date last accessed).
64. Li W, Lin Z, Ashiq MI. *et al.* RoVista: measuring and analyzing the Route Origin Validation (ROV) in RPKI. In: *Internet Measurement Conference (IMC)*. Montreal QC, Canada: Association for Computing Machinery (ACM), 2023.
65. Du B, Testart C, Fontugne R. *et al.* Poster: Taking the Low Road: How RPKI Invalids Propagate. In: *Proceedings of the ACM SIGCOMM 2023 Conference. ACM SIGCOMM '23*. New York, NY: Association for Computing Machinery, 2023, 1144–6.
66. Houston G. ISP Column - March 2021. 2021. https://www.potaroo.net/ispcol/2021-03/rov.html (1 September 2023, date last accessed).
67. APNIC. RPKI I-ROV Filtering World Map. 2021. https://stats.labs.apnic.net/rpki (1 September 2023, date last accessed).
68. Testart C, Richter P, King A. *et al.* To filter or not to filter: measuring the benefits of registering in the RPKI today. In: Sperotto A, Dainotti A, Stiller A. (eds.), *Passive and Active Measurement*. Vol. **12048**, Eugene, Oregon, USA: Springer, 2020.
69. Internet Society. Mutually Agreed Norms for Routing Security (MANRS). https://www.manrs.org/ (1 September 2023, date last accessed).
70. Internet Society. Mutually Agreed Norms for Routing Security (MANRS) Network Operator Participants. https://www.manrs.org/netops/participants/ (1 September 2023, date last accessed).
71. Du B, Testart C, Fontugne R. *et al.* Mind Your MANRS: Measuring the MANRS Ecosystem. In: *ACM Internet Measurement Conference*. Nice, France: Association for Computing Machinery (ACM), 2022.
72. Internet Society. MANRS Network Operator Participants. 2023. https://www.manrs.org/netops/participants/ (1 September 2023, date last accessed).
73. Luckie M, Beverly R, Koga R. *et al.* Network hygiene, incentives, and regulation: deployment of source address validation in the Internet. In: *ACM Computer and Communications Security (CCS)*. London, United Kingdom: Association for Computing Machinery (ACM), 2019.
74. Azimov A, Bogomazov E, Bush R. *et al.* BGP AS_PATH Verification Based on Resource Public Key Infrastructure (RPKI) Autonomous System Provider Authorization (ASPA) Objects. Reston, Virginia, USA: Internet Engineering Task Force, 2023.
75. Salamatian L, Arnold T, Cunha I. *et al.* Who Squats IPv4 Addresses?. *SIGCOMM Comput Commun Rev* 2023;**53**:48–72. https://doi.org/10.1145/3594255.3594260.
76. OECD. BGP incidents, mitigation techniques and policy actions. 2022. https://www.oecd-ilibrary.org/science-and-technology/routing-security_40be69c8-en (1 September 2023, date last accessed).
77. Security and Stability Advisory Committee. SSAC Briefing on Routing Security. 2022. https://www.icann.org/en/system/files/files/sac-121-en.pdf (1 September 2023, date last accessed).
78. Broadband Internet Technical Advisory Group. Security of the Internet's Routing Infrastructure. 2022. https://www.bitag.org/documents/BITAG_Routing_Security.pdf (1 September 2023, date last accessed).
79. Yoo C, Wishnick D. Lowering Legal Barriers to RPKI Adoption. *Faculty Scholarship at Penn Law* 2019. https://scholarship.law.upenn.edu/faculty_scholarship/2035/.
80. American Registry of Internet Numbers. Response to NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing. 2023. https://www.fcc.gov/ecfs/document/10120103512712/1 (1 September 2023, date last accessed).
81. CAIDA. AS relationships data for May 2023. https://publicdata.caida.org/datasets/as-relationships/serial-1/ (1 September 2023, date last accessed).
82. Sriram K, Azimov A. Methods for detection and mitigation of BGP route leaks. In: *Internet Engineering Task Force*. 2022. https://datatracker.ietf.org/doc/draft-ietf-grow-route-leak-detection-mitigation/ (1 September 2023, date last accessed).

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

83. Luckie M, Huffaker B, Dhamdhere A. *et al*. AS relationships, customer cones, and validation. In: *ACM IMC*. Barcelona, Spain: Association for Computing Machinery (ACM), 2013, 243–56.
84. Huffaker B, Luckie M, claffy kc. CAIDA Autonomous System Rankings ASRank. https://asrank.caida.org/ (1 September 2023, date last accessed).
85. CAIDA. CAIDA Internet eXchange Points (IXPs) Dataset. 2023. https://www.caida.org/catalog/datasets/ixps/ (1 September 2023, date last accessed).
86. Routeviews Project–University of Oregon. http://www.routeviews.org/.2022 (1 September 2023, date last accessed).
87. RIPE Routing Information Service. http://www.ripe.net/ris/ (1 September 2023, date last accessed).
88. Job Snijders. NANOG. 2016. https://archive.nanog.org/sites/default/files/Snijders_Everyday_Practical_Bgp.pdf (1 September 2023, date last accessed).
89. Rose S, Borchert O, Mitchell S. *et al*. *Zero Trust Architecture*. Gaithersburg, Maryland, USA: National Institute of Standards and Technology, 2020. https://www.nist.gov/publications/zero-trust-architecture.
90. Council of European Union. Council regulation (EU) no 2016/679, General Data Protection Regulation. 2014. https://eur-lex.europa.eu/eli/reg/2016/679/oj (1 September 2023, date last accessed).
91. Council of European Union. Council regulation (EU) no 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act). 2022. https://eur-lex.europa.eu/eli/reg/2022/2065/oj (1 September 2023, date last accessed).
92. Council of European Union. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on horizontal cybersecurity requirements for products with digital elements and amending Regulation (EU) 2019/1020. 2022. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52022PC0454 (1 September 2023, date last accessed).
93. O'Donovan B. INEX's Shiny New Route Servers. 2019. https://www.inex.ie/inex-news/shiny-new-route-servers/ (1 September 2023, date last accessed).

Downloaded from https://academic.oup.com/cybersecurity/article/10/1/tyae023/7924069 by Georgia Institute of Technology user on 11 February 2025

© The Author(s) 2024. Published by Oxford University Press. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs licence (https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits non-commercial reproduction and distribution of the work, in any medium, provided the original work is not altered or transformed in any way, and that the work is properly cited. For commercial re-use, please contact journals.permissions@oup.com